8

Superimposed Mental Imagery

On the Uses of Make-Perceive

Robert Eamon Briscoe

By absence this good means I gain,
That I can catch her,
Where none can watch her,
In some close corner of my brain…

From 'Present in Absence', attributed to John Donne

1. Introduction: Augmenting Reality with Mental Imagery

The currently most developed and influential theory of visual mental imagery is based on a model of visual object recognition (Kosslyn 1994, 2005; Kosslyn et al. 2006). According to the model, objects are identified when inputs from certain topographically organized areas in the occipital lobe, collectively referred to as the 'visual buffer', are successfully matched against representations stored in long-term memory. When bottom-up inputs from the visual buffer do not clearly specify the presence of a particular object (or kind of object), representations of the features of the best matching object are accessed by an information shunting subsystem. This subsystem performs two strategic, top-down functions. First, it relays information to other subsystems, enabling them to allocate attention to the presumed location of a diagnostic part or feature of the best matching object. Second, it primes the representation of that part or feature in an object-properties processing subsystem to facilitate its encoding. Conscious mental imagery is generated when a stored representation is primed so strongly in the latter system that its activation is propagated backwards along recurrent pathways, inducing a representation of the relevant part or feature in the visual buffer.

Mental images generated during perception, according to this theory, can be used to augment degraded or incomplete perceptual inputs, for example, to enhance representations of a tree partly enveloped in mist or a man hidden in shadow (Kosslyn and

Sussman 1995; Lewis et al. 2011; see also Brockmole et al. 2002). Imagery can plausibly be projected or 'superimposed' on locations in a perceived scene for a variety of other purposes, however. When moving into a new home, for example, one might look at the front door, while simultaneously imagining how a sofa would need to be rotated in order to fit through it. Or, when planning the interior decoration, one might inspect an unfurnished room, while visualizing a bookcase in an empty corner or a carpet on the floor. Yet another familiar example of imaginatively augmented perception is the experience of noticing a constellation in the night-time sky. Noticing a constellation is a hybrid, visual-imaginative experience: it involves both seeing the stars in the constellation and imagining the lines that connect them at the same time. In what follows, I shall refer to such hybrid experiences—involving both a bottom-up, perceptual component and a top-down, imaginative component—as 'make-perceive' (Briscoe 2008, 2011).

In many cases, make-perceive is deliberate and agent initiated: mental images are actively projected and altered in response to incoming visual information and changing task demands. In consequence, the kind of presence that attaches to the imaginatively represented object in cases of active make-perceive, as Brian O'Shaughnessy writes, is ' "thin" and unconvincing' (2000: 349), that is, not liable to being mistaken for real presence. Other cases of make-perceive, however, are passive in that they neither involve effortful visualization nor are susceptible to top-down control. Hence, they involve what O'Shaughnessy calls a 'disturbance of one's sense of reality' (2000: 352). They can sometimes be mistaken for a 'true seeing of a real presence' (2000: 354). This is arguably what is amiss in at least some cases of hallucination: the subject visually imagines an object as present in an otherwise veridically perceived scene, but goes wrong in identifying the internally generated component of her experience as an instance of actually seeing the imagined object.[1]

Passive make-perceive may help to explain two additional phenomena. First, it may account, in part, for illusions in which one's experience of a grey-scale picture (Hansen et al. 2006; Witzel et al. 2011) is influenced by stored information about the depicted object's characteristic colour (this is sometimes referred to as the 'memory-colour effect'). In experiments by Hansen and colleagues, subjects were instructed to adjust to a colour photograph of a banana until it appeared completely achromatic. The picture was generally perceived to be colourless, however, only when its colour was shifted away from neutral grey toward a slightly bluish hue, that is, in a direction opposite to a banana's typical colour. This suggests that when the picture was photometrically achromatic, subjects still perceived it as having a slightly yellow appearance. One empirically plausible explanation of this effect is that viewing the achromatic picture

[1] Passive make-perceive is phenomenologically similar to cases of 'projector' grapheme-colour synaesthesia (Dixon et al. 2004). Projectors report the colours that they experience 'as being "out there on the page", as though a transparency bearing a colored number was placed on top of the written digit' (Dixon et al. 2004: 336). 'Associators', by contrast, describe their colours as 'in the head' or as before their 'mind's eye' (for discussion, see Ward et al. 2007; Mattingley 2009).

elicits a mental image of a yellow banana, which is then amalgamated with bottom-up perceptual signals (Macpherson 2012).[2]

Second, passive make-perceive may also help to explain the finding that subjects sometimes report seeing objects (faces, letters, etc.) in a white-noise display when provided reason to believe that they are present there (Gosselin and Schyns 2003). 'As white noise does not represent coherent structures in the image plane', Gosselin and Schyns write, 'the superstitious perception of a signal had to arise from the observer's share' (2003: 505). Such 'superstitious perception', as they call it, seems well explained by the hypothesis that subjects' expectations caused them to superimpose, unknowingly, a mental image on the screen in front of them.[3]

Importantly, when the observer's share in make-perceive is deliberate or active, the imaginative and perceptual components of the agent's hybrid experience are introspectively distinguishable: the imaginative component, unlike the perceptual component, is subject to top-down control and, so, can be altered or extinguished at will. When make-perceive is passive, however, the imaginative and perceptual components of the agent's experience may be difficult or impossible to tease apart through introspection. One may seem to be in one (purely perceptual) state rather than two (perceptual and imaginative). Classic experiments by Perky (1910) seem to show that agents can under certain conditions mistake what they are seeing for what they are imagining. But cases of passive make-perceive suggest that it is also possible for agents to mistake what they are imagining for what they are seeing.

The examples discussed so far have all involved *visual* perceptual-imaginative hybrids. Make-perceive, however, can cross modal lines: visual perceptual experiences can combine with imagery in non-visual modalities. Seeing a picture of a rose or of a piece of Roquefort cheese, for instance, may elicit concurrent olfactory imagery. Seeing a prickly cactus or a cashmere sweater may elicit tactile or kinaesthetic images of what it would feel like to brush against its surface with one's hand. Seeing a speaker's lips moving in the absence of audible speech sounds may elicit auditory images of the words she is articulating (Calvert et al. 1997; Spence and Deroy 2013).

The perceptual or 'bottom-up' element in make-perceive, it should also be emphasized, need not be visual. When one explores or feels around for an object in a completely dark room at night, one's tactile experiences may be supplemented with projected visual images of the room's layout or the shapes of the various pieces of furniture it contains. In this case, the perceptual contribution is tactile, while the top-down, imaginative contribution comes from vision. (For discussion of touch-driven visual

[2] For a sceptical assessment of the experiments by Hansen et al. (2006) and Witzel et al. (2011), however, see Firestone and Scholl (2015).

[3] The view that seeing and understanding low-information pictures involves the exercise of projective mental imagery is central to the account of pictorial experience developed by E. H. Gombrich in *Art and Illusion*: 'The deliberately blurred image, the *sfumato*, or veiled form . . . cuts down the information on a canvas and thereby stimulates the mechanism of projection' (Gombrich 1961/2000: 175–6). See Briscoe (2018) for discussion.

mental imagery, see Sathian and Zangaladze 2001; Zhang et al. 2004; Lacey and Sathian 2013.) In congenitally blind subjects who make use of echolocation or auditory-tactile sensory substitution devices, bottom-up auditory information about the way objects and surfaces are arrayed in space may be augmented with tactile imagery generated by the 'mind's hand' (Renzi et al. 2013). Many different cross-modal perceptual-imaginative permutations seem to be possible.[4]

My discussion in the rest of this chapter has two parts. In the first part, I show that make-perceive can enable agents to perform certain actions and engage in various kinds of problem-solving more effectively than bottom-up perceiving or top-down imagining alone. In the second part, I turn to the question of whether make-perceive may help to account for the 'phenomenal presence' of occluded or otherwise hidden features of perceived objects (Sellars 1978/2007; Nanay 2010). I argue that phenomenal presence is not well explained by the hypothesis that hidden features are represented using projected mental images. In defending this position, I point to some important phenomenological and functional differences between the way hidden features are represented respectively in mental imagery and amodal completion.

2. Make-Perceive and Problem Solving

The process of superimposing mental imagery on a visually perceived scene is an example of what Gilles Fauconnier refers to as cognitive 'blending'. Fauconnier writes: a 'blend operates in two input mental spaces to yield a third space, the blend. Partial structure from the input spaces is projected into the blended space, which has emergent structure of its own' (1997: 150). In cases of what I am here calling 'make-perceive', the inputs respectively come from perception and imagination, and the emergent blend is a visual-imaginative composite or hybrid experience. Edwin Hutchins (2005) refers to this particular kind of composite as a 'materially anchored blend' since one source of input to the mix is an external, visually perceived scene.

Perception and imagination-based reasoning are powerful modes of non-conceptual cognition indigenous to the biological brain. The examples surveyed in this section serve briefly to illustrate how their blended use enables human agents to solve different types of problems and to carry out certain projects more effectively than by using bottom-up perceiving or top-down imagining alone.[5]

[4] For broad discussion of cross-modal mental imagery 'in which the presentation of a stimulus in one sensory modality results in the formation of a mental image in another modality', see Spence and Deroy (2013).

[5] There are good reasons to think that abilities to 'comment' on visually perceived scenes with auditory imagery (internally recited utterances) play an important role in early childhood word learning and various forms of skill-acquisition (Vygotsky 1962/1986; Diaz and Berk 1992; Dennett 1993; Gauker 2011). Since the cognitive dividends of this kind of visual-auditory make-perceive have been extensively examined elsewhere (see especially Clark 1998), I shall not discuss them here.

2.1 Action guidance

Before deciding what to do or how to move in relation to a perceived scene, human beings can form covert, tactile-kinaesthetic 'motor images' of different possible actions (Jeannerod 2006). For example, before attempting to transport a heavy and unwieldy object, we might imagine different ways of lifting it so as to determine which set of grasp points would minimize torque forces or maximize 'end-state comfort' given what we intend to do with it (Rosenbaum et al. 2006). Or, when engaged in rock climbing, we might imagine different ways of positioning our hands and feet on the surfaces in front of us so as to determine the next set of advantageous holds.[6] In both cases, we not only overlay motor images of possible actions on the visually perceived world, the specific motor images we form are *guided* by incoming perceptual information about the spatial and material properties of objects around us. In this respect, among others, merely imagining the performance of an action in relation to a perceived object utilizes some of the same inner mechanisms as are utilized in the programming stage of overtly executing the action (for a review of empirical findings, see Jeannerod 2006: ch. 2). The ability to engage in such sensorimotor make-perceive—and to thereby anticipate the tactile and kinaesthetic consequences of spatially directed bodily movement—is clearly adaptive: it permits us to simulate and evaluate possible actions 'offline' before risking them overtly in the harsh world (Grush 2004; Vaughan and Zuluaga 2006).

We can covertly imagine moving in relation to the perceived environment. But overt behaviour can also be guided by items and features that we make-perceive in nearby space. Actions performed in processes of artistic and technological creation, for example, frequently depend on the ability imaginatively to add or subtract structure from visually perceived objects. The sculptor (or early tool maker) looks at a piece of stone, visualizes how it would appear if *this* bit were chiselled or knapped away, performs the proper sculpting action, and then evaluates the outcome. This sequence is then reiterated with imaginative modifications to currently existing visual structure determining the goal of action at each step. A similar pattern no doubt characterizes aspects of the design process in drawing (but for limitations on the role of imagery here, see Van Leeuwen et al. 1999).

Other relevant cases involve the production and interpretation of various kinds of pretend behaviour (Van Leeuwen 2011). A child may swing a sword that she visualizes in her empty hand while another child ducks to avoid it. Or a mime may reach for a glass of wine that she and those watching her performance imaginatively project on a nearby table. In cases like these, the motor actions the agent performs are guided by the internally represented spatial, material, and/or functional properties of the objects with which she imaginatively populates space around her.

[6] Thanks to Yoichi Ishida for discussion of this example.

2.2 Diagrammatic reasoning

Numerous studies implicate make-perceive in the interpretation of static machine diagrams. Schwartz and Black (1996), for example, presented subjects with a computer display of two touching gears (Figure 1). Their task was to determine as quickly and accurately as possible whether a knob on one gear and a groove on the other would mesh when the gears were rotated inward. If subjects imagined the rotation of the gears in order to solve this problem, then, it was predicted that their response times would be longer when the knob and groove were placed further apart from the meshing point. And this was just what the experimenters found.

Other findings support the idea that 'mental animation', as Mary Hegarty calls it, can be used to infer the kinematics of a mechanical system from a static visual display. Figure 2 depicts two pulleys. When the free end of the rope is pulled, will the lowermost pulley turn clockwise or counter-clockwise? Studies by Hegarty and colleagues (Hegarty 1992, 2004; Hegarty et al. 2003) have found that when subjects solve problems like this one, they mentally animate the motions of the system's components in a sequence that corresponds to the causal order of visual events in the system's operation.[7] Eye-tracking data are consistent with this account. When asked to predict the motion of a particular component, e.g. the middle pulley in Figure 2, subjects look at that component as well as components earlier in the mechanical process, e.g. the uppermost rope and pulley, but not at components later in the process. The input spaces to the cognitive blend here include the machine diagram and superimposed movement imagery. 'Running the blend' by means of mental animation enables agents to substitute fast, analogue simulations of simple physical interactions for slower, propositionally articulated forms of inference-making.

It should be emphasized that experimental findings such as these are likely to have numerous counterparts in real-world, causal reasoning. As Christopher Gauker (2011)

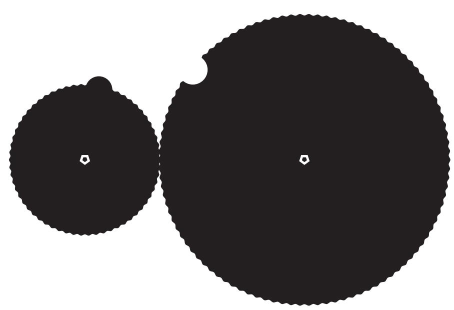

Figure 1. Gears problem. Reproduced with permission from Schwartz and Black (1996).

[7] Hegarty and Steinhoff (1997) and Hegarty and Kozhevnikov (1999) find that mental animation ability is highly related to spatial ability, but not verbal ability. Correspondingly, Sims and Hegarty (1997) find that mental animation interferes more with visuospatial working memory than with verbal working memory.

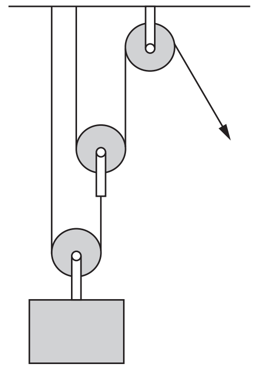

Figure 2. Pulley problem. Reproduced with permission from Hegarty (2004).

has recently argued, many practical problems can be solved using non-conceptual, imaginative representations of how things go together and causally interact. When assembling a piece of furniture from IKEA for example, we may play a game of mental Tetris, trying out possible imaginative fits between the parts spread out on the floor.

2.3 Navigation

My final and most sophisticated example derives from Edwin Hutchins' important studies of long-distance, non-instrumental navigation among the Caroline islanders of Micronesia (Hutchins 1995, 2005). The Caroline islanders, like many other seafaring communities, have learned to use the night-time sky as a compass while at sea. At any given latitude, a star always rises at the same azimuth on the eastern horizon and always sets at same azimuth on the western horizon. A linear constellation or 'star path' is a set of stars that describe the same stationary arc from east to west. Figure 3 illustrates the rising positions of ten of the fourteen linear constellations familiar to Micronesian navigators, with east being the direction of the path for the star Altair. When the bearings of the rising and setting positions are combined, the result is a stable compass in the night-time sky. A skilled navigator can construct the entire compass in imagination from sightings of only one or two stars near the horizon.

The Micronesian sidereal compass performs two main functions at sea. First, it enables the pilot to maintain accurate bearings of distant islands that are well out of sight. From any given point of origin, the pilot knows the star bearing in the direction of which he or she must travel to reach any other island within sailing range. Second, it enables the navigator to keep track of how much of a trip has been completed while travelling on a sealane between two islands. In order to perform this function, however, it must be combined with an impressive application of make-perceive. For every voyage from one island to another, the pilot imagines a third island to the side of the

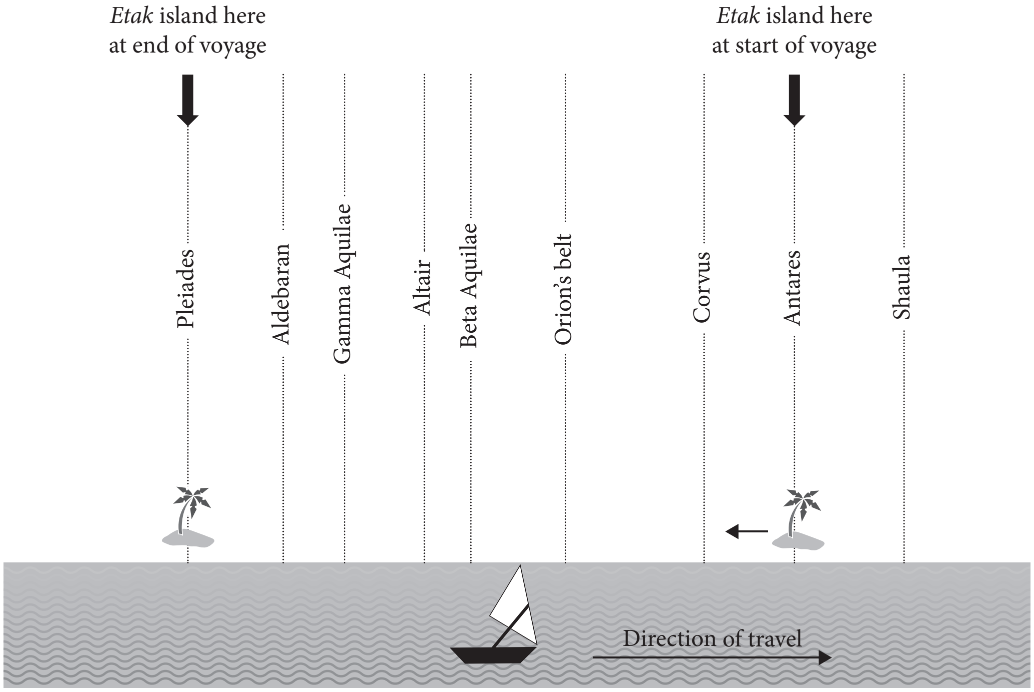

Figure 3. Navigating with reference to the star bearing of an *etak* island. Adapted with permission from Hutchins (2005: 1568).

course and over the horizon, called the *etak* island (Hutchins 2005: 1567–9). Unlike most modern, technologically equipped navigators, Micronesian navigators do not conceive of the voyage using a geocentric spatial framework, that is, in terms of the movement of canoe between two fixed locations on the Earth's surface, but rather egocentrically, in terms of the changing star bearing of the *etak* island relative to the canoe. Hutchins writes:

> at the beginning of the voyage, the *etak* island will be at the star bearing of the *etak* island from the origin. At the end of the voyage, when the canoe has reached the destination, the *etak* island will be at the bearing of the *etak* island from the destination. Thus, during the voyage, the *etak* island appears in the navigator's imagination to move back along the horizon . . . The *etak* island is under one star at the beginning of the voyage and under another at the end of the voyage. By superimposing the imagined movement of the *etak* island on the frame of the star bearings, the Micronesian navigator creates a model of the voyage that he can see and manipulate from his point of view on the deck of the canoe. (2005: 1567–8)

Although space does not permit detailed discussion, it is clear that the *etak* navigation system enables the skilled Micronesian pilot to discern spatial relations and to make inferences in ways that would otherwise be difficult or impossible without the use of maps, tables, GPS, or any of the other external technologies standardly employed by modern seafarers. This example again illustrates the point that strategically combining perception and imagination can yield representational dividends that far surpass their respective, independent contributions to cognition and action planning.

3. Make-Perceive and the Problem of Phenomenal Presence

In section 2, I surveyed some of the ways in which projected or 'materially anchored' mental imagery can facilitate problem solving. I now turn to the question of whether make-perceive may play a much more pervasive and basic role in our everyday visual experience of objects in space around us.

Visual perception is inherently perspectival. One consequence is that from any given position in relation to an opaque, solid object, we only see part of the object's surface: the side of the object that faces us hides its back-side from sight. Another consequence is that objects that are closer in depth often partially occlude those that are further away. Despite these limitations, when observers see an object, they usually have a sense of its presence as a complete, three-dimensional whole.[8] As Nakayama and colleagues write: 'Often we see multiple surfaces in local regions of visual space, with closer objects at least partially covering those behind . . . Yet remarkably, we do not feel much loss of information when part of a surface is rendered invisible by occlusion; we do not see invisible surface regions as nonexistent' (1995: 2). When we see a cat standing behind a picket fence, for example, we see what appears to be a single, intact animal partially hidden by a series of vertical slats. The visible parts of the cat are not experienced as spatially disconnected, but as continuing behind the pickets and as belonging to the same object.

What, however, does it mean to say that 'we do not see invisible surface regions as nonexistent'? How can it be the case that what is hidden from sight is nonetheless experienced as present in the perceived scene? In what follows, I shall refer to this as the problem of 'phenomenal presence'.

In a recent discussion, Bence Nanay treats the problem of phenomenal presence as distinct from the problem of explaining how we represent the hidden features of perceived objects (Nanay 2010). Whereas the former is construed as a phenomenological problem—'[H]ow', Nanay asks, 'can we explain that what it is like to be aware of the occluded parts of perceived objects is similar to what it is like to perceive those parts that are not occluded?' (2010: 252)—the latter is construed as a problem about *representational format*. Are the representations that complete (or, as vision scientists say, 'interpolate') the hidden parts of the cat's body properly perceptual in nature? Or are they rather non-perceptual beliefs that we infer partly on the basis of what we see and partly on the basis of background knowledge?

There is a venerable Kantian–Sellarsian view of the role of imagination in perception that suggests yet a third possibility. According to Sellars, the capacity that Kant (1787/1997) calls productive imagination constructs hybrid 'sense-image models' of

[8] Including young human infants and, perhaps, some non-human animals. For a review of the literature on perceptual completion in human infants, see Bower (1982), Condry et al. (2001), and Otsuka et al. (2006); in monkeys, apes, rodents, and birds, see Mascalzoni and Regolin (2010).

the objects that we perceive. It performs this function, in part, by supplementing awareness of an object's 'occurrent sensible features' with mental images of its hidden features. Sellars writes: 'We do not see *of* the apple its opposite side, or its inside, or its internal whiteness . . . But while these features are not *seen*, they are not *merely* believed in. These features are present in the object of perception as actualities. They are present by virtue of being imagined' (1978/2007: 458).

Nanay adopts this Kantian–Sellarsian view. Like Sellars, he argues that we represent the hidden features of the objects that we perceive by means of projected mental imagery (Nanay 2010: 250). Further, he argues that this view provides, as a corollary, an account of phenomenal presence:

> if what it is like to have visual imagery is similar to what it is like to perceive and being aware of occluded parts of perceived objects is having visual imagery, then, putting these two claims together, we get that what it is like to be aware of the occluded parts of perceived objects is similar to what it is like to perceive those parts that are not occluded. Thus, my proposal that we represent the occluded parts of perceived objects by means of mental imagery has the additional advantage that it gives a simple answer to the question of perceptual presence.
> (2010: 252)

I agree with Sellars and Nanay that we sometimes represent the hidden features of perceived objects by means of projected mental imagery, that is, by means of the capacity that I am calling 'make-perceive'. There are reasons to think, however, that this is far from the whole story. In what follows, I argue, first, that we should distinguish traditional cases of 'amodal completion' (Michotte et al. 1964/1991; Kanizsa 1979; Kanizsa and Gerbino 1982; Wagemans et al. 2012), in which early visual processing mechanisms complete hidden object features on the basis of incoming sensory input, from cases in which the agent generates mental images of hidden object features on the basis of information stored in long-term memory. Second, I argue, contrary to Sellars and Nanay, that the problem of phenomenal presence isn't adequately solved by what I shall call the 'image-based completion' account. The representations that causally support our sense that certain features of perceived objects are really present, though hidden from sight, are properly perceptual representations formed by the mechanisms of amodal completion.

4. Amodal Completion

A standard way of drawing the distinction between 'modal' and 'amodal' completion by students of perception is as follows. In modal completion, the observer characteristically has a distinct, quasi-visual impression of a contour or surface where there are no corresponding stimulus features in the retinal image. Natural scenes that most commonly give rise to modal completion are those in which a foreground surface is camouflaged by a more distant background surface. As a familiar example of this type of completion, consider the illusory Kanizsa squares in Figure 4. Most observers

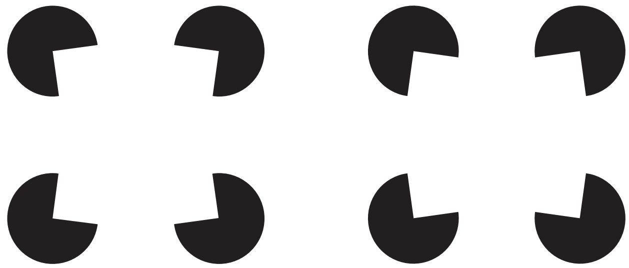

Figure 4. Illusory Kanizsa squares.

report having the impression of seeing a 'thin' square on the left, in which the illusory, vertical contours bow inwards, and the impression of seeing a 'fat' square on the right, in which the illusory, vertical contours bow outwards. The interpolated illusory squares exemplify the 'phenomenal filling-in' characteristic of modal completion (Pessoa et al. 1998).

By contrast, amodal completion occurs when one object is (or appears to be) partially occluded by another and does not typically result in a quasi-visual impression of the object's hidden features. Amodal completion is not characterized by phenomenal filling-in—hence, the epithet 'amodal'. Rather, the phenomenally most salient characteristic of amodal completion is the perceived *unity* of the partially occluded object (Michotte et al. 1964/1991; Kanizsa 1979). When we see a cat walking behind a picket fence, we do not see a moving array of spatially disconnected cat fragments. Rather we see what appears to be a single, intact cat that is partially visible and partially out of sight.

Figure 4, in addition to providing an example of modal completion, also illustrates the phenomenon of amodal completion: the modally completed thin and fat squares appear partially to occlude four black discs. Most subjects have a compelling sense of the continuity of the discs' contours and surfaces *behind* the squares. As another demonstration of amodal completion, consider Figure 5a. Notice here that although it would be reasonable to infer that the occluded object is an octagon given the surrounding context, the completion that the visual system 'prefers' is shown in Figure 5b. This example nicely illustrates the point that the interpolation process in amodal completion 'follows complex principles of its own' (Pylyshyn 1999: 345) and is not rationally sensitive to the observer's beliefs and other high-level cognitive states (for discussion, see Shimojo 2011).

Such *non-cognitive* characterization of amodal completion is supported by empirical evidence that amodally completed contours are represented by stimulus-driven cell-activations in early visual processing areas such as V1 (primary visual cortex) and V2 (Sugita 1999; Bakin et al. 2000; Kamitani and Shimojo 2004; von der Heydt 2004; Komatsu 2006). Sugita (1999), for example, found that amodal completion in area V1 is modulated by binocular disparity. Orientation-selective cells in V1 were presented with two vertical line segments separated by a grey patch. When the patch was presented

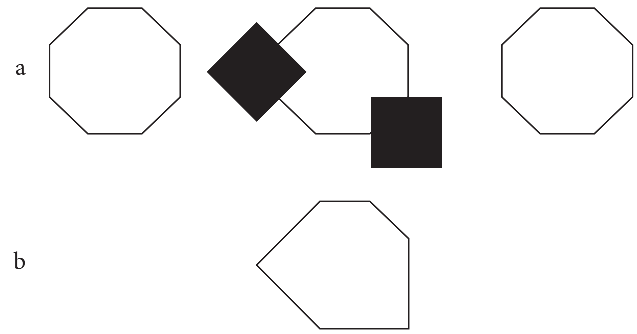

Figure 5. Amodal completion follows its own organizational principles.

with zero disparity or uncrossed disparity, so that it appeared, respectively, on the same or a more distant plane of depth than the line segments, the cells did not respond. However, when the patch was presented with crossed disparity, so that it appeared to be closer than the line segments—a stimulus suggesting the occlusion of a single, vertical bar—the cells responded vigorously. Neuropsychological evidence for rapid completion of occluded objects at early levels of human visual processing is provided by Rensink and Enns (1998) and Johnson and Olshausen (2005). The latter team of investigators found that ERP (event-related potential) differences between images of occluded objects, for example a violin partly hidden by a disc, and images in which object regions are deleted rather than occluded, for example a violin with a disc-shaped cut-out, occur as early as 130 ms after presentation.

There is evidence, it should be emphasized, that the mechanisms of amodal completion extend into higher reaches of the visual processing hierarchy. Hegdé et al. (2008), using fMRI, discovered foci in the lateral occipital complex (area LOC) and the dorsal intraparietal region that are preferentially responsive to partially occluded objects, that is, their response to the presentation of a partially occluded object is significantly larger than their response to either the object or the occluder by itself. In addition, Wokke et al. (2013) in a recent transcranial magnetic stimulation (TMS) study found that feedback to V1/V2 from the LOC plays an important role in perceptual completion in the Kanizsa square illusion.

These findings do not challenge characterization of amodal completion as an essentially non-cognitive, perceptual process. While a small number of subregions within area LOC are preferentially responsive to certain high-level kinds of objects (Grill-Spector and Malach 2004), fMRI studies collectively suggest that LOC, as Nancy Kanwisher puts it, 'exhibits little selectivity for specific object categories' and is largely dedicated to the general-purpose processing of 2D and 3D shapes (Kanwisher 2004: 1184). The findings reported by Hegdé et al. (2008) and Wokke et al. (2013) are thus consistent with the view that amodal completion, while dependent on 'top-down/bottom-up and local-global interactions in a specifically neuroanatomical sense' (Shimojo 2011: 153) are nonetheless early visual processes (Pylyshyn 1999; 2003) and, so, largely independent of cognitively accessible, object-specific knowledge stored in long-term memory.

Before proceeding, two further points are in order. First, important work by Peter Tse (1999) suggests that amodal completion centrally operates at the level of volumes or 3D enclosures. Tse presents a large number of demonstrations that cannot be explained by familiar contour-relatability (Kellman and Shipley 1991) or surface-completion theories (Nakayama and Shimojo 1992; Nakayama et al. 1995), but that are adequately accounted for by his volume-based account. According to Tse, the inputs to the completion process are local surfaces plus the voluminous 'insides' specified by them, and the outputs are maximally closed surfaces in which the local insides are merged. One important implication of this account is that amodal completion is not limited to cases in which an object is partially hidden by an object closer in depth, but will also occur in cases of self-occlusion:

> Amodal completion does not only happen behind an occluder. It is a universal aspect of volume completion, since all objects self-occlude their far side and therefore occlude their true extent. The real problem . . . is not amodal completion at all, but 3-D shape formation or volume completion. What has traditionally been called 'amodal completion' is just a small subset of all volume completion phenomena. (Tse 1999: 62–3)

On this approach, the kind of surface-based completion familiar from demonstrations consisting of flat, overlapping shapes (as in Figures 4–8) is treated as a special case in which completion takes place among ' "degenerate" volumes that do not have insides' (Tse 1999: 42). In section 5, I suggest that our sense of the phenomenal presence of an object's self-occluded parts is plausibly explained by the operation of perceptual mechanisms involved in such non-imaginative, amodal volume completion.

Second, although representations of occluded object-features formed in early visual processing areas do not result in the 'filling-in' characteristic of modal completion, they nonetheless make an important contribution to both the content and phenomenology of conscious visual experience. In particular, they play a significant role in the spatial organization of the 3D scene that we perceive. This point can be brought out by reflecting on Edgar Rubin's 'Maltese Cross' reproduced in Figure 6. The perception evoked by the Maltese Cross is ambiguous or multi-stable, meaning that, with prolonged viewing, figure/ground assignments can alternate. On assignment (a), we visually experience an upright, dark grey cross on a partially occluded white square (completed amodally in the background). On assignment (b), by contrast, we experience a white cross, tilted on its side, partially occluding a dark grey diamond. On a yet third assignment, (c), we experience a grey and white diamond, partially occluding a white square. Notice that it is not only the relative depth relations that flip between these three assignments: there also changes at the level of the *objects* represented in the content of our visual experience. For example, in assignment (b) we experience a white cross and a grey diamond, objects that are absent in assignment (c).

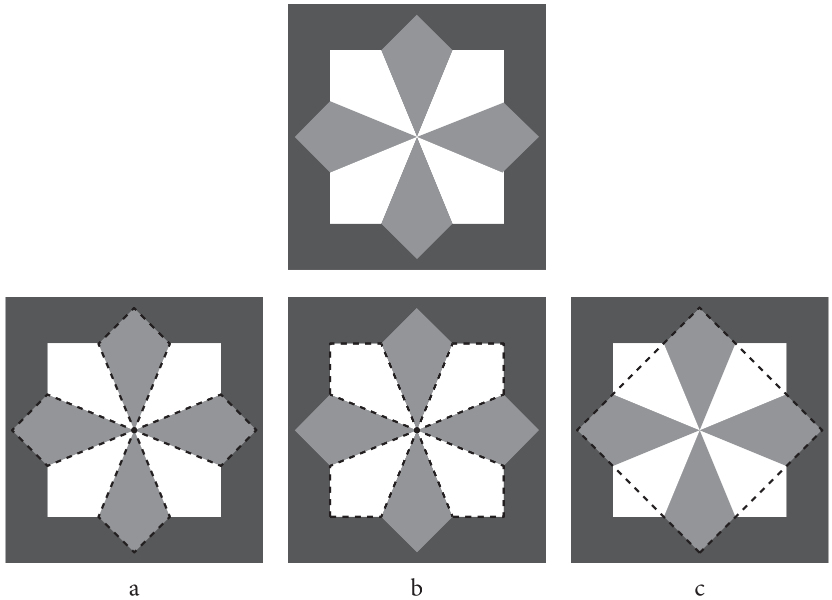

Figure 6. Edgar Rubin's Maltese Cross.

This example nicely illustrates the point that amodally completed contours, surfaces, and volumes are not extra-perceptual *addenda* to what we 'strictly speaking' see. Rather, amodal completion plays an integral role in perceptually segmenting the visual scene into discrete, 3D objects at different distances in depth (Nakayama et al. 1995; Fleming and Anderson 2004; von der Heydt 2004). Indeed, without the stimulus-driven processes that result in amodal completion (and modal completion), the question of how we represent the occluded features of discrete, 3D objects would not arise because we would not *see* coherently organized scenes consisting of such objects (Briscoe 2008, 2011).

5. Image-Based Completion and the Problem of Phenomenal Presence

Not all completion of hidden structure is of a perceptual, non-cognitive character. Completion can also sometimes take place at the level of projected mental imagery or make-perceive.[9] For example, when we see the tip of a pink snout protruding from behind a barn door, we may be disposed to imagine a pig hidden on the other side. Unlike amodal completion, such image-based completion is highly dependent on background knowledge. In particular, it involves identification of the partially

[9] I here focus on visual perceptual completion. See Spence and Deroy (2013) for discussion of some possible roles for non-visual imagery in cross-modal perceptual completion.

occluded object and, hence, accessing categorical information stored in long-term memory. Unlike amodal completion, image-based completion processes are thus cognitively penetrable—they can be 'altered in a way that bears some logical relation to what the person knows' (Pylyshyn 1999: 343). How we imagine the occluded pig attached to the snout depends on our beliefs and memories concerning a pig's normal visual appearance.

The image-based completion account comprises two claims: First, as Nanay writes, '[w]hen we represent the occluded parts of perceived objects, we use mental imagery . . . in a way that would allow us to localize the imagined object in our egocentric space' (2010: 250), that is, we engage in what I am calling make-perceive. Second, the occluded parts of a perceived object are *phenomenally present* in our experience of the object because they are represented using conscious mental imagery.

The basic problem with the first claim is straightforward: it fails to distinguish between non-cognitive, amodal completion and cognitive, image-based completion. Much completion of hidden structure is properly perceptual in nature and does not involve any top-down mental imagery.

Significant phenomenological and functional differences between the two forms of completion, however, also challenge the second claim. First, interpolated contours, surfaces, and volumes in amodal completion (as opposed to modal completion) do not have a visual or quasi-visual phenomenology. What it is like to be aware of the occluded parts of a cat standing behind a picket fence isn't similar to what it is like to be aware of the parts of the cat that are plainly in sight. The occluded parts are phenomenally present—as evidenced in the perceived unity of the cat to which they belong—but they are, as Kanizsa and Gerbino (1982) put it, 'amodally present'. Conscious mental imagery, by contrast, has a modally visual phenomenological character: objects represented in mental imagery are experienced as having certain visible shapes, sizes, and colours. This suggests that conscious mental imagery is not operative in paradigmatic cases of phenomenal presence.

Second, mental images are not typically stable in the absence of sustained effort and fade rapidly. As Hume puts it, 'in the imagination the perception is faint and languid, and cannot without difficulty be preserv'd by the mind steady and uniform for any considerable time' (1739/2000: 11). By contrast, amodal completion phenomena normally persist so long as one perceives their inducers. No more effort is required to experience the partially occluded discs in Figure 4 than to experience the (modally completed) squares that appear to occlude them. In short, image-based completion, unlike amodal completion, is typically introspectively dependent on the agent's own activity. In consequence, its products are not likely to be experienced as having real presence in the perceived scene.

Third, mental images are not *obligatory*. When we see a pink snout protruding from behind a barn door, we may imagine the shape of a hidden pig, but we may also imagine any of a variety of other things instead or indeed nothing at all. By contrast, our experience of amodal completion is not similarly subject to volition or top-down influence. Amodal completion, far from exhibiting what Hume called the 'liberty of the imagination' (1739/2000: 12), operates automatically in accordance with a fairly strict set of organizational principles and is largely driven by bottom-up, sensory inputs. Our experience of amodal completion, in consequence, is not pliant in the way that we would expect it to be were the image-based account correct.

Last, on the Sellars–Nanay account, make-perceive is supposed to explain how we represent hidden features both in cases of superposition, in which an object's visible surfaces hide parts of the more distant background, and *self-occlusion*, in which an object's near side hides its far side. In cases of superposition, both the object's visible surfaces and hidden background regions are represented from a single, unified visual perspective. Visual perception and visual imagination share a common, egocentric point of view. In cases of self-occlusion, however, the perspective of perception and the perspective of imagination come apart. The spatial point of view from which I see the visible surfaces of a car (and egocentrically locate them relative to myself) and the point of view from which I imagine the car's self-occluded surfaces—that is, the surfaces I would see were I counterfactually to view the car from a position facing its far side—are *different* points of view. Hence, it seems unlikely that the Sellars–Nanay account of superposition can be unproblematically extended to cases of self-occlusion.

By contrast, stimulus-driven amodal completion does not only occur in cases of superposition. Studies conducted by Peter Tse, mentioned above in section 4, suggest that the mechanisms of amodal completion typically interpolate self-occluded, volumetric structure.[10] Hence there is reason to suppose that the phenomenal presence of such structure may be causally explained by the construction of representations that are properly perceptual, i.e. non-imaginative, in nature.

6. The Functional Effects of Amodal Completion

At this point in the argument, the following reply might be made on behalf of the image-based account. It is plausible that *active* make-perceive does not account for the phenomenal presence of hidden object features. The products of deliberate imagining or visualization, unlike those of perception, are subject to top-down control and, in consequence, are not experienced as objectively 'out there' in the world (Sartre 1940/2004; Wittgenstein 1953; McGinn 2004). The products of *passive* make-perceive, however, are not so easily distinguished from those of perception. In some cases, they are stable and cannot be altered at will by the agent, e.g. the spots of blood that Lady Macbeth hallucinates on her hands. The same would go for the memory-colour effect,

[10] Although as Tse emphasizes, volume completion processes do not always interpolate the precise form of hidden regions: 'To the extent that an occluded form is interpolated, it may be probabilistic in nature' (Tse 1999: 50).

if, as Macpherson (2012) argues, the underlying mechanism involves 'unbidden' mental imagery. The yellowish appearance of the grey banana picture in the experiments performed by Hansen and colleagues (2006), unlike a product of active, agent-initiated imagining, is stable, automatic, and inducer-specific.

A proponent of image-based completion might thus argue that, when we represent hidden object features, we do so by means of conscious mental images that are passively formed and that are insusceptible to top-down control. The features of objects represented in such imagery would plausibly have phenomenal presence even if those represented in actively projected mental images do not. As Brian O'Shaughnessy writes: 'to the extent that a visual imagining is insightfully experienced as *imagining*, to that extent it is experienced as a mere *quasi-seeing* of an "*unreal presence*", while to the extent that it is *not* so insightfully experienced as imagining, to that same extent it is experienced a true seeing of a real presence' (2000: 354).

While this proposal is not prima facie implausible, it also faces a number of objections. First, as emphasized above, conscious visual mental images have a modally visual phenomenology. The passive projection of such images thus cannot explain the phenomenal presence of object features that are *not* experienced in a modally visual way. What explains the modal phenomenal presence of Duncan's blood on Lady Macbeth's hands cannot explain the amodal phenomenal presence of the hidden parts of the cat behind the picket fence or the discs in Figure 4.

Second, the passive make-perceive account also fails to address the self-occlusion objection voiced in section 5. A solution to the problem of phenomenal presence, however, should plausibly explain both cases of superposition and cases of self-occlusion.

The most serious objection, however, is perhaps the third. As argued above, amodal completion performs a fundamental role in perceptual organization, that is, in grouping and segmenting visible contours, surfaces, and volumes in depth. The amodally completed regions of an object or scene can contribute as much to its perceived spatial organization as do its visible, non-occluded regions. And it is this, no doubt, that confers on the products of amodal completion their distinctive type of phenomenal presence: they have genuinely perceptual functional effects (Kanizsa and Gerbino 1982; Ekroll et al. 2016).

Projected mental images are not suited to play such an organizational functional role. The main reason is that mental image formation, whether active or passive, is sensitive to information stored in long-term visual memory. Perceptual organization, however, is highly insensitive to such top-down influence. Consider, for example, Figure 7. Placing vertical strips on top of the circles displayed on the left side of the figure has a dramatic effect on the spatial organization of the scene we experience: even when we *know* that there is a group of discrete circles behind the strips, we experience the strips as partly occluding a group of continuous vertical waves. Now consider Figure 8. Here we visually experience what appears to be a deer with an extremely long mid-section behind by a grey, rectangular occluder (a 'wiener-dog deer' as my four-year-old son remarked). This perceived organization, however, is at

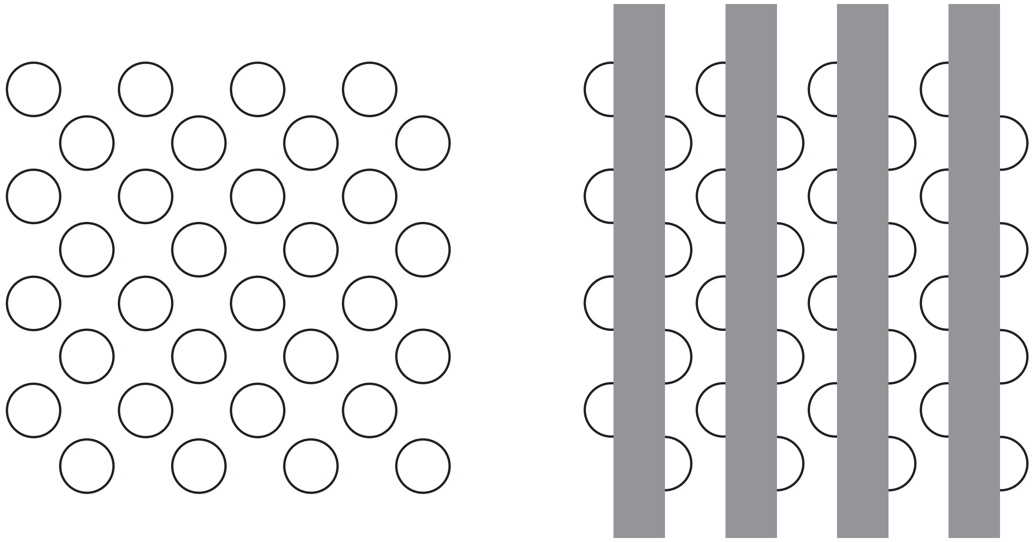

Figure 7. Perceptual organization is sometimes at variance with background knowledge. Reproduced with permission from Ben-Shahar and Ben-Yosef (2015).

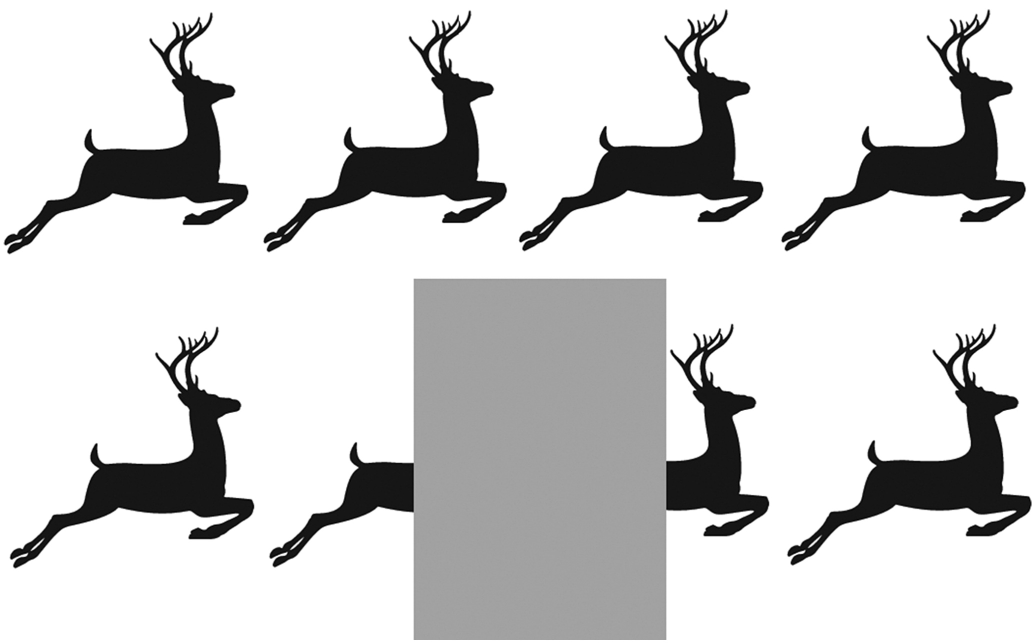

Figure 8. Two partially occluded deer or one elongated deer? Reproduced with permission from Ben-Yosef and Ben-Shahar (2012).

variance with our stored information about deer and their visual appearance properties. It is again not the organization that would be predicted by the image-based completion account.[11]

[11] For numerous other examples of perceptual completions that are at variance with background knowledge, see Kanizsa (1979) and (1985).

7. Conclusion

Imagination, according to Kant, is the faculty by means of which what is perceptually absent is made *present* (1787/1997: B151). It is thus natural to suppose, as do both Sellars and Nanay, that imagination is the faculty by means of which we typically represent the occluded or otherwise hidden features of perceived objects. I have argued here that this natural supposition is false. The representations that give rise to the sense of the phenomenal presence are properly perceptual representations formed on the basis of incoming sensory input from the environment. These representations play a fundamental role in the process of perceptually organizing the scenes we visually experience, a role that conscious mental images are unsuited to play.

As paradoxical as it may sound, we thus sometimes perceive the invisible, that is, we perceive features and parts of objects that do not reflect (or emit) any light to the eye. We do not merely imagine or infer the parts of a cat that are hidden behind the slats in a picket fence. I have tried to show in the foregoing discussion that this view is amply support by empirical work in perceptual psychology. A more developed treatment would survey further sources of evidence and supporting models in different theoretical traditions. Recently influential Bayesian accounts in vision science, for example, model perception as a process of probabilistic inference from the retinal image to its distal source in the environment (Knill and Richards 1996; Kersten et al. 2004; Clark 2013, 2015; Hohwy 2013; Rescorla 2016). A main feature of the Bayesian approach is that the visual system is assumed to have various forms of learned or hardwired 'prior knowledge' about natural scene statistics and the image formation process. Hence, it makes good sense to suppose on this approach that perceptual representations of occluded or obscured contours, surfaces, and volumes will be formed when the bottom-up sensory signal—as interpreted in the light of such prior knowledge—makes their presence probable (Mamassian 2006; Singh and Fulvio 2007; Geisler and Perry 2009). To the extent that there is any kind of puzzle about our perception of the partially occluded cat for Bayesians, Andy Clark suggests, 'it concerns not "presence-in-absence" but (paradoxically) "absence-in-presence"! The puzzle, that is, is why we do not then *only* experience the cat . . . as whole' (Clark 2012: 762).

An earlier and quite compatible view is familiar from work in the Gestalt tradition. 'Unlike their sensationist predecessors', Kellman and Shipley point out, 'the Gestaltists recognized that stimulus variables relevant to perception need not correspond to *local sensations*. Spatial and temporal relationships in the inputs to the senses might explain how perception can instead be in close correspondence to the outside world' (1991: 141, my emphasis).[12] The absence of proximal sensory stimulation caused by a visual

[12] For an explicit defense of the view that Gestalt stimulus factors contributing to amodal completion (such as proximity, good continuation, and symmetry) are probabilistic, ecologically valid indicators of 'life-relevant physical properties of . . . remote environmental objects', see Brunswik and Kimiya (1953).

feature, in other words, does not always entail the absence of stimulus variables that jointly indicate the instantiation of that feature.

Last, from a still influential Gibsonian or 'ecological' perspective, there are a number of sources of visual information for occlusion in the light sampled by the eye, information that 'specifies the existence of one surface behind another, i.e., the continued existence of a hidden surface' (Gibson 1966: 204). These include binocular disparities, T-junctions, and texture accretion/deletion under perspective transformation (for reviews, see Gibson 1979 and Nakayama et al. 1995). For Gibson and contemporary theorists inspired by his work, we sometimes perceive hidden surfaces as directly as we perceive those that hide them.[13]

References

Bakin, Jonathan S., Ken Nakayama, and Charles D. Gilbert (2000). 'Visual responses in monkey areas V1 and V2 to three-dimensional surface configurations', *Journal of Neuroscience* 20: 8188–98.

Ben-Shahar, Ohad and Guy Ben-Yosef (2015). 'Tangent bundle elastica and computer vision', *IEEE Transactions on Pattern Analysis and Machine Intelligence* (*PAMI*) 37: 161–74.

Ben-Yosef, Guy and Ohad Ben-Shahar (2012). 'Tangent bundle curve completion with locally connected parallel networks', *Neural Computation* 24: 3277–316.

Bower, T. G. R. (1982). *Development in Infancy*, 2nd edition. San Francisco, CA: W. H. Freeman.

Briscoe, Robert (2008). 'Vision, action, and make-perceive', *Mind & Language* 23: 457–97.

Briscoe, Robert (2011). 'Mental imagery and the varieties of amodal perception', *Pacific Philosophical Quarterly* 92: 153–73.

Briscoe, Robert (2018). 'Gombrich and the duck-rabbit', to appear in Michael Beaney, Brendan Harrington, and Dominic Shaw (eds.), *Aspect Perception after Wittgenstein: Seeing-As and Novelty*. New York and London: Routledge, 49–88.

Brockmole, James R., Ranxiao F. Wang, and David E. Irwin (2002). 'Temporal integration between visual images and visual percepts', *Journal of Experimental Psychology: Human Perception and Performance* 28: 315–34.

Brunswik, Egon and Joe Kamiya (1953). 'Ecological cue-validity of "proximity" and of other Gestalt factors', *American Journal of Psychology* 66: 20–32.

Calvert, Gemma A., Edward T. Bullmore, Michael J. Brammer, Ruth Campbell, Stephen C. R. Williams, Philip K. McGuire, Peter W. R. Woodruff, Susan D. Iversen, and Anthony S. David (1997). 'Activation of auditory cortex during silent lipreading', *Science* 276: 593–6.

Clark, Andy (1998). 'Magic words: how language augments human computation', in Peter Carruthers and Jill Boucher (eds.), *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press, 162–83.

[13] I am grateful to participants at the 'Perceptual Memory and Perceptual Imagination' conference held at the University of Glasgow in September 2011 for discussion of this chapter, especially Derek Brown, Anya Farennikova, Dominic Gregory, Amy Kind, Fiona Macpherson, and Paul Noordhof. Conversations with Joe Berendzen and Bence Nanay have also helped to shape the ideas presented here.

Clark, Andy (2012). 'Dreaming the whole cat: generative models, predictive processing, and the enactivist conception of perceptual experience', *Mind* 121: 753–71.

Clark, Andy (2013). 'Whatever next? Predictive brains, situated agents, and the future of cognitive science', *Behavioral Brain Sciences* 36: 181–253.

Clark, Andy (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford: Oxford University Press.

Condry, Kirsten F., W. Carter Smith, and Elizabeth F. Spelke (2001). 'Development of perceptual organization', in Francisco Lacerda, Claes von Hofsten, and Mikael Heimann (eds.), *Emerging Cognitive Abilities in Early Infancy*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1–28.

Dennett, Daniel (1993). 'Learning and labeling', *Mind & Language* 8: 540–7.

Diaz, Rafael M. and Laura E. Berk (eds.) (1992). *Private Speech: From Social Interaction to Self-Regulation*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Dixon, Mike J., Daniel Smilek, and Philip M. Merikle (2004). 'Not all synaesthetes are created equal: projector *versus* associator synaesthetes', *Cognitive, Affective, and Behavioral Neuroscience* 4: 335–43.

Ekroll, Vebjørn, Bilge Sayim, Ruth Van der Hallen, and Johan Wagemans (2016). 'Illusory visual completion of an object's invisible backside can make your finger feel shorter', *Current Biology* 26: 1029–33.

Fauconnier, Gilles (1997). *Mappings in Thought and Language*. Cambridge: Cambridge University Press.

Firestone, Chaz and Brian J. Scholl (2015). 'Cognition does not affect perception: evaluating the evidence for "top-down" effects', *Behavioral and Brain Sciences* 39: 1–77.

Fleming, Roland and Barton L. Anderson (2004). 'The perceptual organization of depth', in Leo M. Chalupa and John S. Werner (eds.), *The Visual Neurosciences*, vol. 2. Cambridge, MA: MIT Press, 1284–99.

Gauker, Christopher (2011). *Words and Images: An Essay on the Origin of Ideas*. Oxford: Oxford University Press.

Geisler, Wilson S. and Jeffrey S. Perry (2009). 'Contour statistics in natural images: grouping across occlusions', *Visual Neuroscience* 26: 109–21.

Gibson, James J. (1966). *The Senses Considered as Perceptual Systems*. Boston: Houghton Mifflin.

Gibson, James J. (1979). *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gombrich, E. H. (1961/2000). *Art and Illusion*, 2nd revised edition. Princeton, NJ: Princeton University Press.

Gosselin, Frédéric and Philippe G. Schyns (2003). 'Superstitious perceptions reveal properties of internal representations', *Psychological Science* 14: 505–9.

Grill-Spector, Kalanit and Rafael Malach (2004). 'The human visual cortex', *Annual Review of Neuroscience* 27: 649–77.

Grush, Rick (2004). 'The emulation theory of representation: motor control, imagery, and perception', *Behavioral and Brain Sciences* 27: 377–96.

Hansen, Thorsten, Maria Olkkonen, Sebastian Walter, and Karl R. Gegenfurtner (2006). 'Memory modulates color appearance', *Nature Neuroscience* 9: 1367–8.

Hegarty, Mary (1992). 'Mental animation: inferring motion from static displays of mechanical systems', *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18: 1084–102.

Hegarty, Mary (2004). 'Mechanical reasoning by mental simulation', *Trends in Cognitive Sciences* 8: 280–5.

Hegarty, Mary and Maria Kozhevnikov (1999). 'Types of visual-spatial representations and mathematical problem solving', *Journal of Educational Psychology* 91: 684–9.

Hegarty, Mary, Sarah Kriz, and Christina Cate (2003). 'The roles of mental animations and external animations in understanding mechanical systems', *Cognition and Instruction* 21: 325–60.

Hegarty, Mary and Valerie K. Sims (1994). 'Individual differences in mental animation during mechanical reasoning', *Memory & Cognition* 22: 411–30.

Hegarty, Mary and Kathryn Steinhoff (1997). 'Individual differences in use of diagrams as external memory in mechanical reasoning', *Learning and Individual Differences* 9: 19–42.

Hegdé, Jay, Fang Fang, Scott O. Murray, and Daniel Kersten (2008). 'Preferential responses to occluded objects in the human visual cortex', *Journal of Vision* 8: 1–16.

Hohwy, Jakob (2013). *The Predictive Mind*. Oxford: Oxford University Press.

Hume, David (1739/2000). *A Treatise of Human Nature*, ed. David Norton and Mary Norton. Oxford: Oxford University Press.

Hutchins, Edwin (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.

Hutchins, Edwin (2005). 'Material anchors for conceptual blends', *Journal of Pragmatics* 37: 1555–77.

Jeannerod, Marc (2006). *Motor Cognition: What Actions Tell the Self*. Oxford: Oxford University Press.

Johnson, Jeffrey S. and Bruno A. Olshausen (2005). 'The recognition of partially visible natural objects in the presence and absence of their occluders', *Vision Research* 45: 3262–76.

Kamitani, Yukiyasu and Shinsuke Shimojo (2004). 'Global yet early processing of visual surfaces', in Leo M. Chalupa and John S. Werner (eds.), *The Visual Neurosciences*, vol. 2. Cambridge, MA: MIT Press, 1129–38.

Kanizsa, Gaetano (1979). *Organization in Vision.* New York: Praeger.

Kanizsa, Gaetano (1985). 'Seeing and thinking', *Acta Psychologica* 59: 23–33.

Kanizsa, Gaetano and W. Gerbino (1982). 'Amodal completion: seeing or thinking?', in Jacob Beck (ed.), *Organization and Representation in Perception.* Hillsdale, NJ: Lawrence Erlbaum Associates, 167–90.

Kant, Immanuel (1787/1997). *Critique of Pure Reason*, ed. and trans. Paul Guyer and Allen W. Wood. Cambridge: Cambridge University Press.

Kanwisher, Nancy (2004). 'The ventral visual object pathway in humans: evidence from fMRI', in Leo M. Chalupa and John S. Werner (eds.), *The Visual Neurosciences*, vol. 2. Cambridge, MA: MIT Press, 1179–90.

Kellman, Philip J. and Thomas F. Shipley (1991). 'A theory of visual interpolation in object perception', *Cognitive Psychology* 23: 141–221.

Kersten, Daniel, Pascal Mamassian, and Alan Yuille (2004). 'Object perception as Bayesian inference', *Annual Review of Psychology* 55: 271–304.

Knill, David C. and Whitman Richards (eds.) (1996). *Perception as Bayesian Inference.* New York: Cambridge University Press.

Komatsu, Hidehiko (2006). 'The neural mechanisms of perceptual filling-in', *Nature Reviews Neuroscience* 7: 220–31.

Kosslyn, Stephen M. (1994). *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, MA: MIT Press.

Kosslyn, Stephen M. (2005). 'Mental images and the brain', *Cognitive Neuropsychology* 22: 333–47.

Kosslyn, Stephen M. and Amy L. Sussman (1995). 'Roles of imagery in perception: or, there is no such thing as immaculate perception', in Michael S. Gazzaniga (ed.), *The Cognitive Neurosciences*. Cambridge, MA: MIT Press, 1032–45.

Kosslyn, Stephen M., William L. Thompson, and Giorgio Ganis (2006). *The Case for Mental Imagery*. Oxford: Oxford University Press.

Lacey, Simon and K. Sathian (2013). 'Visual imagery in haptic shape perception', in Simon Lacey and Rebecca Lawson (eds.), *Multisensory Imagery*. Dordrecht: Springer, 207–19.

Lewis, Katie J. S., Grégoire Borst, and Stephen M. Kosslyn (2011). 'Integrating visual mental images and visual percepts: new evidence for depictive representations', *Psychological Research* 75: 259–71.

McGinn, Colin (2004). *Mindsight: Image, Dream, Meaning*. Cambridge, MA: Harvard University Press.

Macpherson, Fiona (2012). 'Cognitive penetration of colour experience: rethinking the issue in light of an indirect mechanism', *Philosophy and Phenomenological Research* 84: 24–62.

Mamassian, Pascal (2006). 'Bayesian inference of form and shape', *Progress in Brain Research* 154: 265–70.

Mascalzoni, Elena and Lucia Regolin (2010). 'Animal visual perception', *WIREs Cognitive Science* 2: 106–16.

Mattingley, Jason B. (2009). 'Attention, automaticity, and awareness in synesthesia', *Annals of the New York Academy of Sciences* 1156: 141–67.

Michotte, Albert, Georges Thinès, and Geneviève Crabbé (1964/1991). 'Amodal completion of perceptual structures', in Georges Thinès, Alan Costall, and George Butterworth (eds.), *Michotte's Experimental Phenomenology of Perception*. Hillsdale, NJ: Lawrence Erlbaum Associates, 140–67.

Nakayama, Ken, Zijiang He, and Shinsuke Shimojo (1995). 'Visual surface representation: a critical link between lower-level and higher-level vision', in Stephen M. Kosslyn and Daniel N. Osherson (eds.), *Visual Cognition*. Cambridge, MA: MIT Press, 1–70.

Nakayama, Ken and Shinsuke Shimojo (1992). 'Experiencing and perceiving visual surfaces', *Science* 257: 1357–63.

Nanay, Bence (2010). 'Perception and imagination: amodal perception as mental imagery', *Philosophical Studies* 150: 239–54.

O'Shaughnessy, Brian (2000). *Consciousness and the World.* Oxford: Oxford University Press.

Otsuka, Yumiko, So Kanazawa, and Masami Yamaguchi (2006). 'Development of modal and amodal completion in infants', *Perception* 35: 1251–64.

Perky, C. W. (1910). 'An experimental study of imagination', *American Journal of Psychology* 21: 422–52.

Pessoa, Luiz, Evan Thompson, and Alva Noë (1998). 'Finding out about filling-in: a guide to perceptual completion for visual science and the philosophy of perception', *Behavioral and Brain Sciences* 21: 723–802.

Pylyshyn, Zenon (1999). 'Is vision continuous with cognition? The case for cognitive impenetrability of visual perception', *Behavioral and Brain Sciences* 22: 341–65.

Pylyshyn, Zenon (2003). *Seeing and Visualizing: It's Not What You Think*. Cambridge, MA: MIT Press.

Rensink, Ronald A. and James T. Enns (1998). 'Early completion of occluded objects', *Vision Research* 38: 2489–505.

Renzi, Chiaraslo, Zaira Cattaneo, Tomaso Vecchi, and Cesare Cornoldi (2013). 'Mental imagery and blindness', in Simon Lacey and Rebecca Lawson (eds.), *Multisensory Imagery*. Dordrecht: Springer, 115–30.

Rescorla, Michael (2016). 'Bayesian perceptual psychology', in Mohan Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*. Oxford: Oxford University Press, 694–716.

Rosenbaum, David A., Rajal G. Cohen, Ruud Meulenbroek, and Jonathan Vaughan (2006). 'Plans for grasping objects', in Markus Latash and Francis Lestienne (eds.), *Motor Control and Learning*. New York: Springer, 9–25.

Sartre, Jean-Paul (1940/2004). *The Imaginary: A Phenomenological Psychology of the Imagination*, trans. Jonathan Webber. London: Routledge.

Sathian, K. and A. Zangaladze (2001). 'Feeling with the mind's eye: the role of visual imagery in tactile perception', *Optometry and Vision Science* 78: 276–81.

Schwartz, Daniel L. and John B. Black (1996). 'Analog imagery in mental model reasoning: depictive models', *Cognitive Psychology* 30: 154–219.

Sellars, Wilfrid (1978/2007). 'The role of the imagination in Kant's theory of experience', in *In the Space of Reasons: Selected Essays of Wilfrid Sellars*, ed. Kevin Scharp and Robert B. Brandom. Cambridge, MA: Harvard University Press, 454–68.

Shimojo, Shinsuke (2011). 'Perceptual organization in the visual cortex', in Liliana Albertazzi, Gert van Tonder, and Dhanraj Vishwanath (eds.), *Perception beyond Inference: The Information Content of Visual Processes*. Cambridge, MA: MIT Press, 137–55.

Sims, Valerie K. and Mary Hegarty (1997). 'Mental animation in the visual-spatial sketchpad: evidence from dual-task studies', *Memory and Cognition* 25: 321–32.

Singh, Manish and Jacqueline M. Fulvio (2007). 'Bayesian contour extrapolation: geometric determinants of good continuation', *Vision Research* 47: 783–98.

Spence, Charles and Ophelia Deroy (2013). 'Crossmodal mental imagery', in Simon Lacey and Rebecca Lawson (eds.), *Multisensory Imagery*. Dordrecht: Springer, 157–83.

Sugita, Yoichi (1999). 'Grouping of image fragments in primary visual cortex', *Nature* 401: 269–72.

Tse, Peter (1999). 'Volume completion', *Cognitive Psychology* 39: 37–68.

Van Leeuwen, C., I. Verstijnen, and P. Hekkert (1999). 'Common unconscious dynamics underlie common conscious effects: a case study in the interactive nature of perception and creation', in J. Scott Jordan (ed.), *Modeling Consciousness across the Disciplines*. Lanham, MD: University Press of America, 179–218.

Van Leeuwen, Neil (2011). 'Imagination is where the action is', *Journal of Philosophy* 108: 55–77.

Vaughan, R. and M. Zuluaga (2006). 'Use your illusion: sensorimotor self-simulation allows complex agents to plan with incomplete self-knowledge', in *Proceedings of the International Conference on Simulation of Adaptive Behaviour (SAB)*. Rome.

von der Heydt, Rüdiger (2004). 'Image parsing mechanisms of the visual cortex', in Leo M. Chalupa and John S. Werner (eds.), *The Visual Neurosciences*, vol. 2. Cambridge, MA: MIT Press, 1139–50.

Vygotsky, Lev (1962/1986). *Thought and Language*. Cambridge, MA: MIT Press.

Wagemans, Johan, James H. Elder, Michael Kubovy, Stephen E. Palmer, Mary A. Peterson, Manish Singh, and Rüdiger von der Heydt (2012). 'A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization', *Psychological Bulletin* 138: 1172–217.

Ward, Jamie, Ryan Li, Shireen Salih, and Noam Sagiv (2007). 'Varieties of grapheme-colour synaesthesia: a new theory of phenomenological and behavioural differences', *Consciousness and Cognition* 16: 913–31.

Wittgenstein, Ludwig (1953). *Philosophical Investigations*, trans. G. E. M. Anscombe. Oxford: Blackwell.

Witzel, Christoph, Hanna Valkova, Thorsten Hansen, and Karl R. Gegenfurtner (2011). 'Object knowledge modulates colour appearance', *i-Perception* 2: 13–49.

Wokke, Martijn, Annelinde Vandenbroucke, H. Steven Scholte, and Victor Lamme (2013). 'Confuse your illusion: feedback to early visual cortex contributes to perceptual completion', *Psychological Science* 24: 63–71.

Zhang, Minming, Valerie D. Weisser, Randall Stilla, S. C. Prather, and K. Sathian (2004). 'Multisensory cortical processing of object shape and its relation to mental imagery', *Cognitive, Affective, and Behavioral Neuroscience* 4: 251–9.